arXiv:1507.01101v2 [cs.DC] 22 Oct 2015

Utility Maximizing Thread Assignment and Resource Allocation

Pan Lai, Rui Fan, Wei Zhang, Fang Liu

School of Computer Engineering
Nanyang Technological University, Singapore
{plai1, fanrui}@ntu.edu.sg, {wei.zhang,fliu}@ieee.org

***Abstract*—Achieving high performance in many distributed systems requires finding a good assignment of threads to servers as well as effectively allocating each server's resources to its assigned threads. The assignment and allocation components of this problem have been studied largely separately in the literature. In this paper, we introduce the *assign and allocate (AA)* problem, which seeks to simultaneously find an assignment and allocations that maximize the total utility of the threads. Assigning and allocating the threads together can result in substantially better overall utility than performing the steps separately, as is traditionally done. We model each thread by a concave utility function giving its throughput as a function of its assigned resources. We first show that the AA problem is NP-hard, even when there are only two servers. We then present a $2(\sqrt{2}-1) > 0.828$ factor approximation algorithm, which runs in $O(mn^2 + n(\log mC)^2)$ time for n threads and m servers with C amount of resources each. We also present a faster algorithm with the same approximation ratio and $O(n(\log mC)^2)$ running time. We conducted experiments to test the performance of our algorithm on threads with different types of utility functions, and found that it achieves over 99% of the optimal utility on average. We also compared our algorithm against several other assignment and allocation algorithms, and found that it achieves up to 5.7 times better total utility.**

Keywords: Resource allocation; algorithms; multicores; cloud

I. Introduction

In this paper, we study efficient ways to run a set of threads on multiple servers. Our problem consists of two steps. First, each thread is assigned to a server. Subsequently, the resources at each server are allocated to the threads assigned to it. Each thread obtains a certain utility based on the resources it obtains. The goal is to maximize, over all possible assignments and allocations, the total utility of all the threads. We call this problem *AA*, for *assign and allocate*.

While our terminology generically refers to threads and servers, the AA problem can model a range of distributed systems. For example, in a multicore processor, each core corresponds to a server offering its shared cache as a resource to concurrently executing threads. Each thread is first bound to a core, after which cache partitioning [4], [3] can enforce an allocation of the core's cache among the assigned threads. Each thread's performance is a (typically concave) function of the size of its cache partition. As the threads' cache requirements differ, achieving high overall utility requires both a proper mapping of threads to the cores and properly partitioning each core's cache. Another application of AA is a web hosting center in which a number of web service threads run on multiple servers and compete for resources such as processing or memory. The host seeks to maximize system utility to run a larger number of services and obtain greater revenue. Lastly, consider a cloud computing setting in which a provider sells virtual machine instances (threads) running on physical machines (servers). Customers express their willingness to pay for instances consuming different amounts of resources using utility functions, and the provider's task is to assign and size the virtual machines to maximize her profit.

The two steps in AA correspond to the *thread assignment* and *resource allocation* problems, both of which have been studied extensively in the literature. However, to our knowledge, these problems have not been studied together in the unified context considered in this paper. Existing works on resource allocation [12], [16], [4], [3], [11] largely deal with dividing the resources on a single server among a given set of threads. It is not clear how to apply these algorithms when there are multiple servers, since there are many possible ways to initially assign the threads to servers, and certain assignments result in low overall performance regardless of how resources are subsequently allocated. For example, if there are two types of threads, one with high maximum utility and one with low utility, then assigning all the high utility threads to the same server will result in competition between them and low overall utility no matter how resources are allocated. Likewise, existing works on thread assignment [7], [8], [5], [6] often overlook the resource allocation aspect. Typically each thread requests a fixed amount of resource, and once assigned to a server it is allocated precisely the resource it requested,

without any adjustments based on the requests of the other threads assigned to the same server. This can also lead to suboptimal performance. For example, consider a thread that obtains x^β utility when allocated x amount of resource, for some $\beta \in (0, 1)$[1], and suppose the thread requests z resources, for an arbitrary z. Then when there are n threads and one server with C resources, $\frac{C}{z}$ of the threads would receive z resources each while the rest receive 0, leading to a total utility of $Cz^{\beta-1}$, which is constant in n. However, the optimal allocation gives $\frac{C}{n}$ resources to each thread and has total utility $C^\beta n^{1-\beta}$, which is arbitrarily better for large n.

The AA problem combines thread assignment and resource allocation. We model each thread using a non-decreasing and concave *utility function* giving its performance as a function of the resources it receives. The goal is to simultaneously find assignments and allocations for all the threads that maximizes their total utility. We first show that the AA problem is NP-hard, even when there are only two servers. In contrast, the problem is efficiently solvable when there is a single server [16]. Next, we present an approximation algorithm that finds a solution with utility at least $\alpha = 2(\sqrt{2}-1) > 0.828$ times the optimal. The algorithm relates the optimal solution of a single server problem to an approximately optimal solution of the multiple server problem. The algorithm runs in time $O(mn^2 + n(\log mC)^2)$, where n and m are the number of threads and servers, respectively, and C is the amount of resource on each server. We then present a faster algorithm with $O(n(\log mC)^2)$ running time and the same approximation ratio. Finally, we conduct experiments to test the performance of our algorithm. We used several types of synthetic threads with different utility functions, and show that our algorithm obtains at least 99% of the maximum utility on average for all the thread types. We also compared our algorithm with several simple but practical assignment and allocation heuristics, and show that our algorithm performs up to 5.7 times better when the threads have widely differing utility functions.

The rest of paper is organized as follows. In Section 2, we describe related works on thread assignment and resource allocation. Sections 3 formally defines our model and the AA problem. Section 4 proves AA is NP-hard. Section 5 presents an approximation algorithm and its analysis, and Section 6 proposes a faster algorithm. In Section 7, we describe our experimental results. Finally, Section 8 concludes.

II. Related Works

There is a large body of work on resource allocation for a single server. Fox *et al.* [12] considered concave utility functions and proposed a greedy algorithm to find an optimal allocation in $O(nC)$ time, where n is the number of threads and C is the amount of resource on the server. Galil [16] proposed an improved algorithm with $O(n(\log C)^2)$ running time, by doing a binary search to find an allocation in which the derivatives of all the threads' utility functions are equal, and the total resources used by the allocation is C. Resource allocation for nonconcave utility functions is weakly NP-complete. However, Lai and Fan [11] identified a structural property of real-world utility functions that leads to fast parameterized algorithms. In a multiprocessor setting, utility typically corresponds to cache miss rate. Miss rate curves can be determined by running threads multiple times using different cache allocations. Qureshi *et al.* [4] proposed efficient hardware based methods to minimize the overhead of this process, and also designed algorithms to partition a shared cache between multiple threads using methods based on [12]. Zhao *et al.* [10] proposed software based methods to measure page misses and allocate memory to virtual machines in a multicore system. Chase *et al.* [9] used utility functions to dynamically determine resource allocations among users in hosting centers and maximize total profit.

There has also been extensive work on assigning threads to cores in multicore architectures. Becchi *et al.* [7] proposed a scheme to determine an optimized assignment of threads to heterogeneous cores. They characterized the behavior of each thread on a core by a single value, its IPC (instructions per clock), which is independent of the amount of resource the thread uses. In contrast, we model a thread using a utility function giving its performance for different resource allocations. Radojković *et al.* [8] proposed a statistical approach to assign threads on massively multithreaded processors, choosing the best assignment out of a large random sample. The results in [7] and [8] are based on simulations and provide no theoretical guarantees. A problem related to thread assignment is *application placement*, in which applications with different resource requirements need to be mapped to servers while fulfilling certain quality of service guarantees. Urgaokar *et al.* [5] proposed offline and online approximation algorithms for application placement. The offline algorithm achieves a $\frac{1}{2}$ approximation ratio. However, these algorithms are not directly comparable to ours, as they consider multiple types of resources while we consider one type.

Co-scheduling is a technique which divides a set of threads into subsets and executes each subset together on one chip of a multicore processor. It is frequently used to minimize cache interference between threads. Jiang et al. [13] proposed algorithms to find optimal co-schedules for pairs of threads, and also approximation algorithms for co-scheduling larger groups of threads.

[1]For $\beta = \frac{1}{2}$, this is known as the "square root rule" [1], [2].

Tian *et al.* [14] also proposed exact and approximation algorithms for co-scheduling on chip multiprocessors with a shared last level cache. Zhuralev *et al.* [15] surveyed scheduling techniques for sharing resources in multicores. The works on co-scheduling require measuring the performance from running different groups of threads together. When co-scheduling large groups, the number of measurements required becomes prohibitive. In contrast, we model threads by utility functions, which can be determined by measuring the performance of individual threads instead of groups.

The AA problem is related to the multiple knapsack and also multiple-choice knapsack (MCKP) problems. There are many works on both problems. For the former, Neebe *et al.* [20] proposed a branch-and-bound algorithm, and Chekuri *et al.* [21] proposed a PTAS. The multiple knapsack problem differs from AA in that each item, corresponding to a thread, has a single weight and value, corresponding to a single resource allocation and associated throughput; in contrast, we use utility functions which allow threads a continuous range of allocations and throughputs. The MCKP problem can model utility functions as it considers classes of items with different weights and values and chooses one item from each class; the class corresponds to a utility function. However, MCKP only considers a single knapsack, and thus corresponds to a restricted form of AA with one server. Kellerer [17] proposed a greedy MCKP algorithm. Lawler [19] proposed a $1-\epsilon$ approximate algorithm, while Gens and Levner [18] proposed a $\frac{4}{5}$ approximate algorithm with better running time. AA can be seen as a combined multiple-choice multiple knapsack problem. We are not aware of any previous work on this problem. This paper is a first step for the case where the ratios of values to weights in each item class is concave and there are items for every weight.

III. Model

In this section we formally define our model and problem. We consider a set of m homogeneous servers $s_1, \ldots, s_m$, where each server has $C > 0$ amount of resources. The homogeneity assumption is reasonable in a number of settings. For example, multicore processors typically contain shared caches of the same size, and datacenters often have many identically configured servers for ease of management. Homogeneous servers have also been widely studied in the literature [5], [23]. We also have n threads $t_1, \ldots, t_n$. We imagine the threads are performing long-running tasks, and so the set of threads is static. Let S and T denote the set of servers and threads, respectively. Each thread t_i is characterized by a *utility function* $f_i : [0, C] \to \mathbb{R}^{\geq 0}$, giving its performance as a function of the resources it receives. We assume that f_i is nonnegative, nondecreasing and concave. The concavity assumption models a diminishing returns property frequently observed in practice [4]. While it does not apply in all settings, concavity is a common assumption used in many utility models, especially on cache and memory performance [1], [2].

Our goal is to assign the threads to the servers in a way that respects the resource bounds and maximizes the total utility. While a solution to this problem involves both an assignment of threads and allocations of resources, for simplicity we use the term assignment to refer to both. Thus, an *assignment* is given by a vector $[(r_1, c_1), \ldots, (r_n, c_n)]$, indicating that each thread t_i is allocated c_i amount of resource on server s_{r_i}. Let S_j be the set of threads assigned to server s_j. That is, $S_j = \{i \,|\, r_i = j\}$. Then for all $1 \leq j \leq m$, we require $\sum_{i \in S_j} c_i \leq C$, so that the threads assigned to s_j use at most C resources. We assume that every thread is assigned to some server, even it receives 0 resources on the server. The *total utility* from an assignment is $\sum_{j=1}^{m} \sum_{i \in S_j} f_i(c_i) = \sum_{i \in T} f_i(c_i)$. The *AA (assign and allocate)* problem is to find an assignment that maximizes the total utility.

IV. Hardness of the Problem

In this section, we show that it is NP-hard to find an assignment maximizing the total utility, even when there are only two servers. Thus, it is unlikely there exists an efficient optimal algorithm for the AA problem. This motivates the approximation algorithms we present in Sections V and VI.

Theorem IV.1. *Finding an optimal AA assignment is NP-hard, even when there are only two servers.*

Proof: We give a reduction from the NP-hard *partition problem* [22] to the AA problem with two servers. In the partition problem, we are given a set of numbers $S = \{c_1, \ldots, c_n\}$ and need to determine if there exists a partition of S into sets S_1 and S_2 such that $\sum_{i \in S_1} c_i = \sum_{i \in S_2} c_i$. Given an instance of partition, we create an AA instance A with two servers each with $C = \frac{1}{2} \sum_{i=1}^{n} c_i$ amount of resources. There are n threads $t_1, \ldots, t_n$, where the i'th thread has utility function f_i defined by

$$f_i(x) = \begin{cases} x & \text{if } x \leq c_i \\ c_i & \text{otherwise} \end{cases}$$

The f_i functions are nondecreasing and concave. We claim the partition instance has a solution if and only if A's maximum utility is $\sum_{i=1}^{n} c_i$. For the if direction, let $A^* = [(r_1^*, c_1^*), \ldots, (r_n^*, c_n^*)]$ denote an optimal solution for A, and let S_1 and S_2 be the set of threads assigned to the servers 1 and 2, respectively. We show that S_1, S_2 solve the partition problem. We first show that $c_i^* = c_i$ for all i. Indeed, if $c_i^* < c_i$ for some i,

then $f_i(c_i^*) < c_i$, while $f_j(c_j^*) \leq c_j$, for all $j \neq i$. Thus, $\sum_{j=1}^n f_j(c_j^*) < \sum_{j=1}^n c_j$, which contradicts the assumption that A^*'s utility is $\sum_{j=1}^n c_j$. Next, suppose $c_i^* > c_i$ for some i. Then since A^* is a valid assignment, we have $\sum_{i \in S_1} c_i^* + \sum_{i \in S_2} c_i^* \leq C + C = \sum_{i=1}^n c_i$, and so there exists $j \neq i$ such that $c_j^* < c_j$. But then $f_j(c_j^*) < c_j$ and $f_k(c_k^*) \leq c_k$ for all $k \neq j$, so A^*'s total utility is $\sum_{k=1}^n f_k(c_k^*) < \sum_{i=1}^n c_i$, again a contradiction. Thus, we have $c_i^* = c_i$ for all i, and so $\sum_{i \in S_1} c_i^* + \sum_{i \in S_2} c_i^* = \sum_{i=1}^n c_i = 2C$. So, since $\sum_{i \in S_1} c_i^* \leq C$ and $\sum_{i \in S_2} c_i^* \leq C$, then $\sum_{i \in S_1} c_i^* = \sum_{i \in S_2} c_i^* = C$. Hence, $\sum_{i \in S_1} c_i = \sum_{i \in S_2} c_i = C$, and S_1, S_2 solve the partition problem.

For the only if direction, suppose S_1, S_2 are a solution to the partition instance. Then since $\sum_{i \in S_1} c_i = \sum_{i \in S_2} c_i = C$, we can assign the threads with indices in S_1 and S_2 to servers 1 and 2, respectively, and get a valid assignment with utility $\sum_{i=1}^n f_i(c_i) = \sum_{i=1}^n c_i$. This is a maximum utility assignment for A, since $f_i(x) \leq c_i$ for all i. Thus, the partition problem reduces to the AA problem, and so the latter is NP-hard for two servers. ■

V. Approximation Algorithm

In this section, we present an algorithm to find an assignment with total utility at least $\alpha = 2(\sqrt{2} - 1) > 0.828$ times the optimal in $O(mn^2 + n(\log mC)^2)$ time. The algorithm consists of two main steps. The first step transforms the utility functions, which are arbitrary nondecreasing concave functions and hence difficult to work with algorithmically, into a function consisting of two linear segments which is easier to handle. Next, we find an α-approximate optimal thread assignment for the linearized functions. We then show that this leads to an α approximate solution for the original concave problem.

A. Linearization

To describe the linearization procedure, we start with the following definition.

Definition V.1. *Given an AA problem A with m servers each with C amount of resources, and n threads with utility functions $f_1, \ldots, f_n$, a* super-optimal allocation *for A is a set of values $\hat{c}_1, \ldots, \hat{c}_n$ that maximizes the quantity $\sum_{i=1}^n f_i(\hat{c}_i)$, subject to $\sum_{i=1}^n \hat{c}_i \leq mC$. Call $\sum_{i=1}^n f_i(\hat{c}_i)$ the* super-optimal utility *for A.*

To motivate the above definition, note that for any valid assignment $[(r_1, c_1), \ldots, (r_n, c_n)]$ for A, we have $\sum_{i=1} c_i \leq mC$, and so the utility of the assignment $\sum_{i=1}^n f_i(c_i)$ is at most the super-optimal utility of A. Let F^* and $\hat{F}$ denote A's maximum utility and super-optimal utility, respectively. Then we have the following.

Lemma V.2. $F^* \leq \hat{F}$.

Thus, to find an α approximate solution to A, it suffices to find an assignment with total utility at least $\alpha\hat{F}$. We note that the problem of finding $\hat{F}$ and the associated super-optimal allocation can be solved in $O(n(\log mC)^2)$ time using the algorithm from [16], since the f_i functions are concave. Also, since these functions are nondecreasing, we have the following.

Lemma V.3. $\sum_{i=1}^n \hat{c}_i = mC$.

In the remainder of this section, fix A to be an AA problem consisting of m threads with C resources each and n threads with utility functions $f_1, \ldots, f_n$. Let $[\hat{c}_1, \ldots, \hat{c}_n]$ be a super-optimal allocation for A computed as in [16]. We define the *linearized* version of A to be another AA problem B with the same set of servers and threads, but where the threads have piecewise linear utility functions $g_1, \ldots, g_n$ defined by

$$g_i(x) = \begin{cases} f_i(\hat{c}_i)\frac{x}{\hat{c}_i} & \text{if } x < \hat{c}_i \\ f_i(\hat{c}_i) & \text{otherwise} \end{cases} \tag{1}$$

Lemma V.4. *For any $i \in T$ and $x \in [0, C]$, $f_i(x) \geq g_i(x)$.*

Proof: For $x \in [0, \hat{c}_i]$, we have $f_i(x) \geq \frac{\hat{c}_i - x}{\hat{c}_i} f_i(0) + \frac{x}{\hat{c}_i} f_i(\hat{c}_i) \geq g_i(x)$, where the first inequality follows because f_i is concave, and the second inequality follows because $f_i(0) \geq 0$. Also, for $x > \hat{c}_i$, $f_i(x) \geq f_i(\hat{c}_i) = g_i(x)$. ■

B. Approximation algorithm for linearized problem

We now describe an α approximation algorithm for the linearized problem. The pseudocode is given in Algorithm 1. The algorithm takes as input a super-optimal allocation $[\hat{c}_1, \ldots, \hat{c}_n]$ for A and the resulting linearized utility functions $g_1, \ldots, g_n$, as described in Section V-A. Variable C_j represents the amount of resource left on server j, and R is the set of unassigned threads. The outer loop of the algorithm runs until all threads in R have been assigned. During each loop, U is the set of (thread, server) pairs such that the server has at least as much remaining resource as the thread's super-optimal allocation. If any such pairs exist, then in line 6 we find a thread in U with the greatest utility using its super-optimal allocation, breaking ties arbitrarily. Otherwise, in line 9 we find a thread that can obtain the greatest utility using the remaining resources of any server. In both cases we assign the thread in line 12 to a server giving it the greatest utility. Lastly, we update the server's remaining resources accordingly.

C. Analyzing the linearized algorithm

We now analyze the quality of the assignment produced by Algorithm 1. We first define some notation. Let $D = \{i \in T \mid c_i = \hat{c}_i\}$ be the set of threads whose allocation in Algorithm 1 equals its super-optimal allocation, and let $E = T - D$ be the remaining threads.

Algorithm 1

Input: Super-optimal allocation $[\hat{c}_1, \ldots, \hat{c}_n]$, and $g_1, \ldots, g_n$ as defined in Equation 1

1: $C_j \leftarrow C$ for $j = 1, \ldots, m$
2: $R \leftarrow \{1, \ldots, n\}$
3: **while** $R \neq \emptyset$ **do**
4: $\quad U \leftarrow \{(i,j) \,|\, (i \in R) \wedge (1 \leq j \leq m) \wedge (\hat{c}_i \leq C_j)\}$
5: $\quad$ **if** $U \neq \emptyset$ **then**
6: $\quad\quad (i,j) \leftarrow \text{argmax}_{(i,j)\in U}\ g_i(\hat{c}_i)$
7: $\quad\quad c_i \leftarrow \hat{c}_i$
8: $\quad$ **else**
9: $\quad\quad (i,j) \leftarrow \text{argmax}_{i\in T, 1\leq j\leq m}\ g_i(C_j)$
10: $\quad\quad c_i \leftarrow C_j$
11: $\quad$ **end if**
12: $\quad r_i \leftarrow j$
13: $\quad R \leftarrow R - \{i\}$
14: $\quad C_j \leftarrow C_j - c_i$
15: **end while**
16: **return** $(r_1, c_1), \ldots, (r_n, c_n)$

We say the threads in D are *full*, and the threads in E are *unfull*. Note that full threads are the ones computed in line 6, and unfull threads are computed in line 9. The full threads have the same utility in the super-optimal allocation and the allocation produced by Algorithm 1. Thus, to show Algorithm 1 achieves a good approximation ratio it suffices to show the utilities of the unfull threads in Algorithm 1 are sufficiently large compared to their utilities in the super-optimal allocation. We first show some basic properties about the unfull threads.

Lemma V.5. *At most one thread from E is assigned to any server.*

Proof: Suppose for contradiction there are two threads $t_a, t_b \in E$ assigned to a server s_k, and assume that t_a was assigned before t_b. Consider the time of t_b's assignment, and let S_j denote the set of threads assigned to a server s_j. We have $\sum_{i\in S_k} c_i = C$, because $t_a \in E$, and so it was allocated all of s_k's remaining resources in lines 10 and 14 of Algorithm 1. Also, $\sum_{i\in S_j} c_i = C$ for any $j \neq k$. Indeed, if $\sum_{i\in S_j} c_i < C$ for any $j \neq k$, then s_j has more remaining resources than s_k, and so t_b would be assigned to s_j instead of s_k because it can obtain more utility. Thus, together we have that when t_b is assigned, $\sum_{i\in T} c_i \geq \sum_{j=1}^{m} \sum_{i\in S_j} c_i = mC$. Now, $c_i \leq \hat{c}_i$ for all $t_i \in T$. Also, since $t_a, t_b \in E$, then $c_a < \hat{c}_a$ and $c_b < \hat{c}_b$. Thus, we have $\sum_{i\in T} \hat{c}_i > \sum_{i\in T} c_i \geq mC$, which is a contradiction because $\sum_{i\in T} \hat{c}_i = mC$ by Lemma V.3. ∎

Lemma V.6. $|E| \leq m - 1$.

Proof: Lemma V.5 implies that $|E| \leq m$, so it suffices to show $|E| \neq m$. Assume for contradiction $|E| = m$. Then by Lemma V.5, for each server s_k there exists a thread $t_a \in E$ assigned to s_k. t_a receives all of s_k's remaining resources, and so $\sum_{i\in S_k} c_i = C$ after its assignment. Then after all m threads in E have been assigned, we have $\sum_{i\in T} c_i = mC$. But since $c_a < \hat{c}_a$ for all $t_a \in E$, and $c_i \leq \hat{c}_i$ for all $t_i \in T$, we have $\sum_{i\in T} \hat{c}_i > \sum_{i\in T} c_i = mC$, which is a contradiction. Thus, $|E| \neq m$ and the lemma follows. ∎

The next lemma shows that the total resources allocated to the unfull threads in Algorithm 1 is not too small compared to their super-optimal allocation.

Lemma V.7. $\sum_{i\in E} c_i \geq \frac{|E|}{m} \sum_{i\in E} \hat{c}_i$.

Proof: We first partition the servers into sets U and V, where $U = \{j \in S \,|\, S_j \subseteq D\}$ is the set of servers containing only full threads, and $V = S - U$ are the servers containing some unfull threads. Let $C_j = C - \sum_{i\in S_j} c_i$ be the amount of unused resources on a server s_j at the end of Algorithm 1. Then $C_j = 0$ for all $j \in V$, since the unfull thread in S_j was allocated all the remaining resources on s_j. So, we have $\sum_{j\in U} C_j = \sum_{j\in S} C_j = \sum_{j\in S}(C - \sum_{i\in S_j} c_i) = mC - \sum_{i\in T} c_i$, and so $\sum_{i\in T} c_i = mC - \sum_{i\in U} C_j$.

Next, we have $\sum_{i\in T} c_i = \sum_{i\in D} \hat{c}_i + \sum_{i\in E} c_i = mC - \sum_{i\in E} \hat{c}_i + \sum_{i\in E} c_i$. The first equality follows because $c_i = \hat{c}_i$ for $i \in D$, and the second equality follows because $D \cup E = T$ and $\sum_{i\in T} \hat{c}_i = mC$. Combining this with the earlier expression for $\sum_{i\in T} c_i$, we have $mC - \sum_{i\in U} C_j = mC - \sum_{i\in E} \hat{c}_i + \sum_{i\in E} c_i$, and so

$$\sum_{i\in E} c_i + \sum_{i\in U} C_i = \sum_{i\in E} \hat{c}_i. \tag{2}$$

Now, assume for contradiction that $\sum_{i\in E} c_i < \frac{|E|}{m} \sum_{i\in E} \hat{c}_i$. Then by Equation 2 we have

$$\sum_{i\in U} C_i > \frac{m - |E|}{m} \sum_{i\in E} \hat{c}_i. \tag{3}$$

We have $|V| = |E|$, since by Lemma V.5 each server in V contains only one unfull thread. Thus $|U| = m - |V| = m - |E|$. Using this in Equation 3, we have that there exists an $j \in U$ with

$$C_j \geq \frac{1}{|U|} \sum_{i\in U} C_i > \frac{1}{m} \sum_{i\in E} \hat{c}_i. \tag{4}$$

We claim that for all $i \in E, j \in U$, $c_i \geq C_j$. Indeed, suppose $c_i < C_j$ for some i. But since $C_j > c_i$, t_i should be allocated to s_j because it can obtain greater utility on s_j than its current server, which is a contradiction. Thus, $c_i \geq C_j$ for all $i \in E$. Using this and Equation 4, we

have

$$\sum_{i \in E} c_i \geq \sum_{i \in E} C_j = |E| C_j > \frac{|E|}{m} \sum_{i \in E} \hat{c}_i$$

However, this contradicts the assumption that $\sum_{i \in E} c_i < \frac{|E|}{m} \sum_{i \in E} \hat{c}_i$. Thus, the lemma follows. ∎

Let $\gamma = \max_{i \in E} g_i(\hat{c}_i)$ be the maximum super-optimal utility of any thread in E. The following lemma says that all of the first m threads assigned by Algorithm 1 are given their super-optimal allocations and have utility at least γ.

Lemma V.8. *Let t_i be one of the first m threads assigned by Algorithm 1. Then $t_i \in D$ and $g_i(c_i) \geq \gamma$.*

Proof: To show $t_i \in D$, note that the m servers all had C resource at the start of Algorithm 1, and fewer than m threads were assigned before t_i. So when t_i was assigned, there was at least one server with C resource. Then t_i can obtain $\hat{c}_i$ resource on one of these servers, and so $t_i \in D$.

To show $g_i(c_i) \geq \gamma$, suppose the opposite, and let $j \in E$ be such that $g_j(\hat{c}_j) = \gamma$. Since $\hat{c}_j \leq C$, and since in t_i's iteration there was some server with C resource, then in that iteration Algorithm 1 would have obtained greater utility by assigning t_j instead of t_i, which is a contradiction. Thus, $g_i(c_i) \geq \gamma$. ∎

Lemma V.8 implies there are at least m threads in D, and so we have the following.

Corollary V.9. $\sum_{i \in D} g_i(c_i) \geq m\gamma$.

The next lemma shows that for the threads in E, threads with higher slopes in the nonconstant portion of their utility functions are allocated more resources.

Lemma V.10. *For any two threads $i, j \in E$, if $\frac{g_i(\hat{c}_i)}{\hat{c}_i} > \frac{g_j(\hat{c}_j)}{\hat{c}_j}$, then $c_i \geq c_j$.*

Proof: Suppose for contradiction $c_i < c_j$, and suppose first that t_i was assigned before t_j. Then when t_i was assigned, there was at least one server with c_j or more remaining resources. We have $\hat{c}_i > c_j$, since otherwise t_i can be allocated $\hat{c}_i$ resources, so that $i \notin E$. Now, since $\hat{c}_i > c_j > c_i$, then t_i could obtain greater utility by being allocated c_j instead of c_i amount of resources. This is a contradiction.

Next, suppose t_j was assigned before t_i. Then when t_j was assigned, there was a server with at least c_j amount of resources. Again, we have $\hat{c}_i > c_j$. Indeed, otherwise we have $\hat{c}_i \leq c_j$, and $\hat{c}_j > c_j$ since $j \in E$, and so t_i can be allocated its super-optimal allocation while t_j cannot. But Algorithm 1 prefers in line 4 to assign threads that can receive their super-optimal allocations, and so it would assign t_i before t_j, a contradiction. Thus, $\hat{c}_i > c_j$. However, this means that in the iteration in which t_j was assigned, t_i can obtain greater utility than t_j, since $g_i(c_j) = c_j \frac{g_i(\hat{c}_i)}{\hat{c}_i} > g_j(c_j) = c_j \frac{g_j(\hat{c}_j)}{\hat{c}_j}$, where the first equality follows because $\hat{c}_i > c_j$, the inequality follows because $\frac{g_i(\hat{c}_i)}{\hat{c}_i} > \frac{g_j(\hat{c}_j)}{\hat{c}_j}$, and the second equality follows because $\hat{c}_j > c_j$. Thus, t_i would be assigned before t_j, a contradiction. The lemma thus follows. ∎

The following facts are used in later parts of the proof. The first fact follows from the Cauchy-Schwarz inequality, and the second is Chebyshev's sum inequality.

Fact V.11. *Given $a_1, \ldots, a_n > 0$, we have $(\sum_{i=1}^n a_i)(\sum_{i=1}^n \frac{1}{a_i}) \geq n^2$.*

Fact V.12. *Given $a, a', b, b' > 0$, if $\frac{a}{a'} < \frac{b}{b'}$, then $\frac{a}{a'} < \frac{a+b}{a'+b'} < \frac{b}{b'}$.*

Fact V.13. *Given $a_1 \geq a_2 \geq \ldots \geq a_n$ and $b_1 \geq b_2 \geq \ldots \geq b_n$, we have $\sum_{i=1}^n a_i b_i \geq (\frac{1}{n} \sum_{i=1}^n a_i)(\sum_{i=1}^n b_i)$.*

We now state a lower bound on a certain function.

Lemma V.14. *Let $A, d > 0$, and $0 < a_1 \leq a_2 \ldots \leq a_n$. Also, let $\beta = (A + \sum_{i=1}^n a_i z_i)/(A + \sum_{i=1}^n z_i)$, where each $z_i \in [0, d]$. Then*

$$\beta \geq \min_{j=1,\ldots,n} \left(\frac{A + \sum_{i=1}^j a_i d}{A + jd}, 1 \right)$$

Proof: If $a_1 \geq 1$, then Fact V.12 implies that $\beta \geq 1$, and the lemma holds. Otherwise, suppose $a_1 < 1$. Then differentiating β with respect to z_1, we get

$$\beta'(z_1) = \frac{(a_1 - 1)A + \sum_{i=2}^n (a_1 - a_i) z_i}{(A + \sum_{i=1}^n z_i)^2}$$

Since $a_1 \leq a_2 \ldots \leq a_n$ and $a_1 < 1$, we have $\beta'(z_1) < 0$. Thus, $\beta(z_1)$ is minimized for $z_1 = d$, and we have

$$\beta \geq \frac{A + a_1 d + \sum_{i=2}^n a_i z_i}{A + d + \sum_{i=2}^j z_i}$$

To simplify this expression, suppose first that $(A + a_1 d)/(A + d) \leq a_2$. Then we have

$$\beta \geq \frac{A + a_1 d + \sum_{i=2}^n a_i z_i}{A + d + \sum_{i=2}^n z_i} \geq \frac{A + a_1 d}{A + d}.$$

The second inequality follows because $a_2 \leq \ldots \leq a_n$ and by Fact V.12. Thus, the lemma is proved. Otherwise, $(A + a_1 d)/(A + d) > a_2$, and so

$$\frac{A + a_1 d + \sum_{i=2}^n a_i z_i}{A + d + \sum_{i=2}^n z_i} \geq \frac{A + a_1 d + a_2 d + \sum_{i=3}^n a_i z_i}{A + 2d + \sum_{i=3}^n z_i}$$

We can simplify the latter expression in a way similar to above, based on whether $(A + a_1 d + a_2 d)/(A + 2d) \leq a_3$. Continuing this way, if we stop at the j'th step, then $\beta \geq (A + \sum_{i=1}^j a_i d)/(A + jd)$. Otherwise, after the n'th step, we have $\beta \geq (A + \sum_{i=1}^n a_i d)/(A + nd)$. In either case, the lemma holds. ∎

Algorithm 1 produces an allocation $c_1, \ldots, c_n$ with total utility $G = \sum_{i \in D} g_i(\hat{c}_i) + \sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i} c_i$. We now

prove that this allocation is an α approximation to the super-optimal utility $\hat{F} = \sum_{i \in T} f_i(\hat{c}_i)$.

Lemma V.15. *$G \geq \alpha \hat{F}$, where $\alpha = 2(\sqrt{2}-1) > 0.828$.*

Proof: We have $\hat{F} = \sum_{i \in T} f_i(\hat{c}_i) = \sum_{i \in T} g_i(\hat{c}_i)$ by the definition of the g_i's. Thus,

$$\begin{aligned}
\frac{G}{\hat{F}} &= \frac{\sum_{i \in D} g_i(\hat{c}_i) + \sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i} c_i}{\sum_{i \in D} g_i(\hat{c}_i) + \sum_{i \in E} g_i(\hat{c}_i)} \\
&\geq \frac{m\gamma + \sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i} c_i}{m\gamma + \sum_{i \in E} g_i(\hat{c}_i)} \\
&\geq \frac{m\gamma + \left(\sum_{i \in E} c_i/|E|\right) \sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i}}{m\gamma + \sum_{i \in E} g_i(\hat{c}_i)} \\
&\geq \frac{m\gamma + \left(\sum_{i \in E} \hat{c}_i/m\right) \sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i}}{m\gamma + \sum_{i \in E} g_i(\hat{c}_i)}
\end{aligned}$$

Recall that $\gamma = \max_{i \in E} g_i(\hat{c}_i)$. The first inequality follows because of Corollary V.9 and because $\hat{c}_i \geq c_i$ for all i. The second inequality follows because by Lemma V.10, threads $i \in E$ with larger values of $\frac{g_i(\hat{c}_i)}{\hat{c}_i}$ also have larger values of c_i. Thus, we can apply Fact V.13 to bring the term $\sum_{i \in E} c_i/|E|$ outside the sum $\sum_{i \in E} \frac{g_i(\hat{c}_i)}{\hat{c}_i} c_i$. The last inequality follows because of Lemma V.7. Now, assume WLOG that the elements in E are ordered by nonincreasing value of $\hat{c}_i$, so that $\frac{1}{\hat{c}_1} \leq \frac{1}{\hat{c}_2} \leq \ldots \leq \frac{1}{\hat{c}_{|E|}}$. Let E_i denote the first i elements of E in this order. For any $i \in E$, we have $g_i(\hat{c}_i) \in [0, \gamma]$. Thus, applying Lemma V.14 to the last expression above, letting $g_i(\hat{c}_i)$ play the role of z_i and $\frac{\sum_{i \in E} \hat{c}_i}{m\hat{c}_i}$ play the role of a_i, and noting $\frac{G}{\hat{F}} \leq 1$, we have

$$\begin{aligned}
\frac{G}{\hat{F}} &\geq \min_{i=1,\ldots,|E|} \left(\frac{m\gamma + \left(\sum_{j \in E} \hat{c}_j/m\right) \sum_{j \in E_i} \frac{\gamma}{\hat{c}_j}}{m\gamma + i\gamma} \right) \\
&\geq \min_{i=1,\ldots,|E|} \left(\frac{m + \frac{1}{m}\left(\sum_{j \in E_i} \hat{c}_j\right)\left(\sum_{j \in E_i} \frac{1}{\hat{c}_j}\right)}{m + i} \right) \\
&\geq \min_{i=1,\ldots,|E|} \left(\frac{m + \frac{i^2}{m}}{m + i} \right)
\end{aligned}$$

The second inequality follows by simplification and because $\sum_{j \in E} \hat{c}_j \geq \sum_{j \in E_i} \hat{c}_j$ for any i. The last inequality follows by Fact V.11. It remains to lower bound the final expression. Treating i as a real value and taking the derivative with respect to i, we find the minimum value is obtained at $i = (\sqrt{2}-1)m$, for which $\frac{G}{\hat{F}} \geq 2(\sqrt{2}-1) = \alpha$. Thus, the lemma is proved. ∎

D. Solving the concave problem

To solve the original AA problem with concave utility functions $f_1, \ldots, f_n$, we run Algorithm 1 on the linearized problem to obtain an allocation $c_1, \ldots, c_n$, then simply output this as the solution to the concave problem. The total utility of this solution is $F = \sum_{i \in T} f_i(c_i)$. We now show this is an α approximation to the optimal utility F^*.

Theorem V.16. *$F \geq \alpha F^*$, and Algorithm 1 achieves an α approximation ratio.*

Proof: We have $F = \sum_{i \in T} f_i(c_i) \geq \sum_{i \in T} g_i(c_i) \geq \alpha \hat{F} \geq \alpha F^*$, where the first inequality follows because $f_i(c_i) \geq g_i(c_i)$ by Lemma V.4, the second inequality follows by Lemma V.15, and the last inequality follows by Lemma V.2. ∎

Next, we give a simple example that shows our analysis of Algorithm 1 is nearly tight.

Theorem V.17. *There exists an instance of AA where Algorithm 1 achieves $\frac{5}{6} > 0.833$ times the optimal total utility.*

Proof: Consider 3 threads, and 2 servers each with one (divisible) unit of resource. Let

$$f_1(x) = \begin{cases} 2x & \text{if } x \in [0, \frac{1}{2}] \\ 1 & \text{if } x > \frac{1}{2} \end{cases}$$

Also, let $f_2(x) = x$. Suppose the first two threads both have utility functions f_1, and third thread has utility function f_2. The super-optimal allocation is $[\hat{c}_1, \hat{c}_2, \hat{c}_3] = [\frac{1}{2}, \frac{1}{2}, 1]$. Algorithm 1 may assign threads 1 and 2 to different servers, with $\frac{1}{2}$ resource each, then assign thread 3 to server 1 with $\frac{1}{2}$ resource. This has a total utility of $2\frac{1}{2}$. On the other hand, the optimal assignment is to put threads 1 and 2 on server 1 and thread 3 on server 2. This has a utility of 3. ∎

Lastly, we analyze Algorithm 1's time complexity.

Theorem V.18. *Algorithm 1 runs in $O(mn^2 + n(\log mC)^2)$ time.*

Proof: Computing the super-optimal allocation takes $O(n(\log mC)^2)$ time using the algorithm in [16]. Then the algorithm runs n loops, where in each loop it computes the set U with $O(mn)$ elements. Thus, the theorem follows. ∎

VI. A Faster Algorithm

In this section, we present a faster approximation algorithm that achieves the same approximation ratio as Algorithm 1 in $O(n(\log mC)^2)$ time. The pseudocode is shown in Algorithm 2.

The algorithm also takes as input a super-optimal allocation $\hat{c}_1, \ldots, \hat{c}_n$, which we compute as in Section V-A. It sorts the threads in nonincreasing order of $g_i(\hat{c}_i)$. It then takes threads $m+1$ to n in this ordering, and resorts them in nonincreasing order of $g_i(\hat{c}_i)/\hat{c}_i$. Next, it initializes $C_1, \ldots, C_m$ to C, and stores them in a

max heap H. C_j represents the amount of remaining resources on server j. The main loop of the algorithm iterates through the threads in order. Each time it chooses the server with the most remaining resources, allocates the minimum of the thread's super-optimal allocation and the server's remaining resources to it, and assigns the thread to the server. Then H is updated accordingly.

Algorithm 2

Input: Super-optimal allocation $[\hat{c}_1, \ldots, \hat{c}_n]$, and $g_1, \ldots, g_n$ as defined in Equation 1

1: sort threads in nonincreasing order of $g_i(\hat{c}_i)$ as $t_1, \ldots, t_n$
2: sort $t_{m+1}, \ldots, t_n$ in nonincreasing order of $g_i(\hat{c}_i)/\hat{c}_i$
3: $C_j \leftarrow C$ for $j = 1, \ldots, m$
4: store $C_1, \ldots, C_m$ in a max-heap H
5: **for** $i = 1, \ldots, n$ **do**
6: $\quad j \leftarrow \text{argmax}_{1 \le j \le m} C_j$
7: $\quad c_i \leftarrow \min(\hat{c}_i, C_j)$
8: $\quad C_j \leftarrow C_j - c_i$, and update H
9: $\quad r_i \leftarrow j$
10: **end for**
11: **return** $(r_1, c_1), \ldots, (r_n, c_n)$

A. Algorithm analysis

We now show Algorithm 2 achieves a $2(\sqrt{2}-1)$ approximation ratio, and runs in $O(n(\log mC)^2)$ time. The proof of the approximation ratio uses exactly the same set of lemmas as in Section V-A and V-B. The proofs for most of the lemmas are also similar. Rather than replicating them, we will go through the lemmas and point out any differences in the proofs. Please refer to Sections V-A and V-B for the definitions, lemma statements and original proofs.

- *Lemmas V.2, V.3, V.4* These lemmas deal with the super-optimal allocation, which is the same in Algorithms 1 and 2.
- *Lemma V.5* The proof of this lemma depended on the fact that in Algorithm 1 if we assign a second E thread to a server, then all the other servers have no remaining resources. This is also true in Algorithm 2, since in line 6 we assign a thread to a server with the most remaining resources, and so if when we assign a second E thread t to a server s there was another server s' with positive remaining resources, we would assign t to s' instead, a contradiction.
- *Lemma V.6* This follows by exactly the same arguments as the original proof.
- *Lemma V.7* The only statement we need to check from the original proof is that for all $i \in E$ we have $c_i \ge C_j$. But this is true in Algorithm 2 because if there were any $c_i < C_j$, line 6 of Algorithm 2 would assign thread i to server j instead of i's current server, a contradiction. All the other statements in the original proof then follow.
- *Lemma V.8* This follows because lines 1 and 2 of Algorithm 2 show that the first m assigned threads have at least as much super-optimal utility as the remaining $n-m$ threads. Also, the first m threads must be in D, since there is always a server with C resources during the first m iterations of Algorithm 2. Thus, all threads in E are among the last $n-m$ assigned threads, and their maximum super-optimal utility is no more than the minimum utility of any D thread.
- *Corollary V.9* This follows immediately from Lemma V.8.
- *Lemma V.10* As we stated above, all threads in E must be among the last $n-m$ assigned by Algorithm 2. That is, they are among threads $t_{m+1}, \ldots, t_n$. In line 2 these threads are sorted in nondecreasing order of $g_i(\hat{c}_i)/\hat{c}_i$. Thus, the lemma follows.
- *Facts V.11 to V.13, Lemma V.14* These follow independently of Algorithm 2.
- *Lemma V.15* The proof of this lemma used only the preceding lemmas, not any properties of Algorithm 1. Thus, it also holds for Algorithm 2.

Given the preceding lemmas, we can state the approximation ratio of Algorithm 2. The proof of the theorem is the same as the proof of Theorem V.16, and is omitted.

Theorem VI.1. *Let F be the total utility from the assignment produced by Algorithm 2, and let F^* be the optimal total utility. Then $F \ge \alpha F^*$.*

Lastly, we analyze Algorithm 2's time complexity.

Theorem VI.2. *Algorithm 2 runs in $O(n(\log mC)^2)$ time.*

Proof: Finding the super-optimal allocation takes $O(n(\log mC)^2$ time using the algorithm in [16]. Steps 1 and 2 take $O(n \log n)$ time. Since C is usually large in practice, we can assume that $\log n = O(\log mC)^2$. Each iteration of the main for loop takes $O(\log m)$ time to extract the maximum element from H and update H. Thus, the entire for loop takes $O(n \log m)$ time. Thus, the overall running time is dominated by the time to find a super-optimal allocation, and the theorem follows. ■

VII. Experimental Evaluation

In this section we evaluate the performance of our algorithms experimentally. As both Algorithms 1 and 2 have the same approximation ratio, we only evaluate Algorithm 2. We compare the total utility the algorithm achieves with the super-optimal (SO) utility, which is at least as large as the optimal utility. We also compare the

algorithm with several simple but practical heuristics we name UU, UR, RU and RR. The UU (uniform-uniform) heuristic assigns threads in a round robin manner to the servers, and allocates the threads assigned to a server the same amount of resources. UR (uniform-random) assigns threads in a round robin manner, and allocates threads a random amount of resources on each server. RU (random-uniform) assigns threads to random servers, and equally allocates resources on each server. Finally, RR (random-random) randomly assigns threads and allocates them random amounts of resource.

Our experiments use threads with random concave utility functions generated according to various probability distributions, as follows. We fix an amount of resources C on each server, and set the value of the utility function at 0 to be 0. Then we generate two values v and w according to the distribution H, conditioned on $w \leq v$, and set the value of the utility function at $\frac{C}{2}$ to v, and the value at C to $v+w$. Lastly, we apply the PCHIP interpolation function from Matlab to the three generated points to produce a smooth concave utility function. In all the experiments, we set the number of servers to be $m = 8$, and test the effects of varying different parameters. One parameter is $\beta = \frac{n}{m}$, the average number of threads per server. All the experiments run quickly in practice. Using $m = 8, n = 100$ and $C = 1000$, an unoptimized Matlab implementation of Algorithm 2 finishes in only 0.02 seconds. The results below show the average performance from 1000 random trials.

A. Uniform and normal distributions

We first consider Algorithm 2's performance compared to SO, UU, UR, RU and RR on threads with utility functions generated according to the uniform and normal distributions. We set the mean and standard deviation of the normal distribution to be 1 and 1, respectively. Figures 1(a) and 1(b) show the ratio of Algorithm 2's total utility versus the utilities of the other algorithms, for β varying between 1 to 15. The behaviors for both distributions are similar. Compared to SO, our performance never drops below 0.99, meaning that Algorithm 2 always achieves at least 99% of the optimal utility. The ratios of our total utility compared to those of UU, UR, RU and RR are always above 1, so we always perform better than the simple heuristics. For small values of β, UU performs well. Indeed, for $\beta = 1$, UU achieves the optimal utility because it places one thread on each server and allocates it all the resources. UR does not achieve optimal utility even for $\beta = 1$, since it allocates threads random amounts of resources. RU and RR may allocate multiple threads per server, and also do not achieve the optimal utility. As β grows, the performance of the heuristics gets worse relative to Algorithm 2. This is because as the number of threads grows, it becomes more likely that some threads have very high maximum utility. These threads need to be assigned and allocated carefully. For example, they should be assigned to different servers and allocated as much resources as possible. The heuristics likely fail to do this, and hence obtain low performance. The performance of UR and RR, as well as those of UU and RU converge as β grows. This is because both random and uniform assignments assign the threads roughly evenly between the servers for large β. Also, the performance of UU and RU is substantially better than UR and RR, which indicates that the way in which resources are allocated has a bigger effect on performance than how threads are assigned, and that uniform allocation is generally better than random allocation.

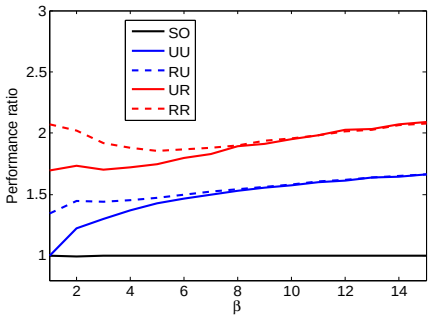

(a) Uniform distribution

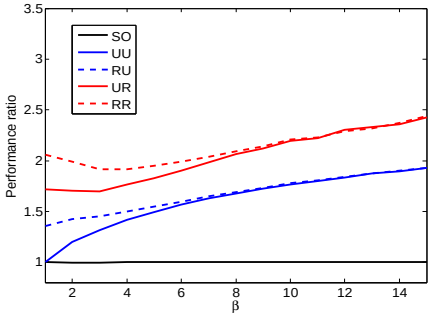

(b) Normal distributionn

Fig. 1. Performance of Algorithm 2 versus SO,UU,RU,UR, and RR as a function of β under the uniform and normal distributions.

B. Power law distribution

We now look at the performance of Algorithm 2 using threads generated according to the power law distribution. Here, each value x has a probability $\lambda x^{-\alpha}$ of occurring, for some $\alpha > 1$ and normalization factor λ. Figure 2(a) shows the effect of varying β while fixing $\alpha = 2$. Here we see the same trends as those under the uniform and normal distributions, namely that Algorithm 2 always performs very close to optimal, while the performance of the heuristics gets worse with increasing β. However, the rate of performance degradation is faster

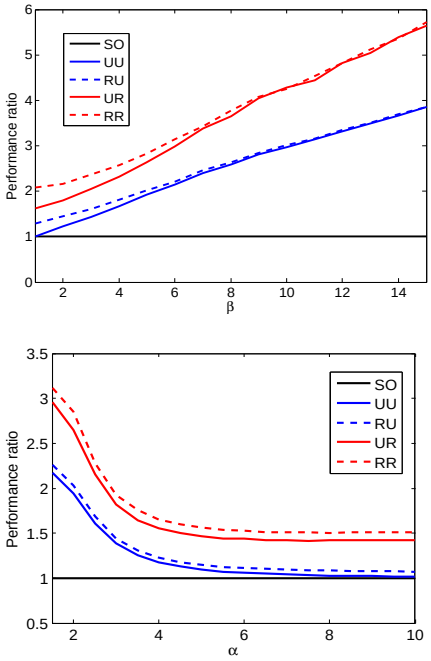

Fig. 2. Performance of Algorithm 2 versus SO,UU,RU,UR, and RR as a function of β and α under the power law distribution.

than with the uniform and normal distributions. This is because the power law distribution with $\alpha = 2$ is more likely to generate threads with very different maximum utilities. These threads must be carefully assigned and allocated, which the heuristics fail to do. For $\beta = 15$, Algorithm 2 is 3.9 times better than UU and RU, and 5.7 times better than UR and RR.

Figure 2(b) shows the effect of varying α, using a fixed $\beta = 5$. Algorithm 2's performance is nearly optimal. In addition, the performance of the heuristics improves as α increases. This is because for higher values of α, it is unlikely that there are threads with very high maximum utilities. So, since the maximum utilties of the threads are roughly the same, almost any even assignment of the threads works well. Despite this, we still observe that UU and RU perform better than UR and RR. This is because when the threads are roughly the same, then due to the concavity of the utility functions, the optimal allocation is to give each thread the same amount of resources. This is done by UU and RU, but not by UR and RR, which allocate resources randomly.

C. Discrete distribution

Lastly, we look at the performance using utility functions generated by a discrete distribution. This distribution takes on only two values ℓ, h, with $\ell < h$. γ is a parameter that controls the probability that ℓ occurs, and $\theta = \frac{h}{\ell}$ is a parameter that controls the relative size of the values. Figure 3(a) shows Algorithm 2's performance as we vary β, fixing $\gamma = 0.85$ and $\theta = 5$. The same trends as with the other distributions are observed. Figure 3(b) shows the effect of varying γ, when $\beta = 5$ and $\theta = 5$. Our algorithm achieves the lowest performance for $\gamma = 0.75$, when we achieve 97.5% of the super-optimal utility. The four heuristics also perform worst for this value. For γ close to 0 or 1, all the heuristics perform well, since these correspond to instances where either h or ℓ is very likely to occur, so that almost all the threads have the same maximum utility. Lastly, we consider the effect of varying θ. Here, as θ increases, the difference between the high and low utilities becomes more evident, and the effects of poor thread assignments or misallocating resources become more serious. Hence, the performance of the heuristics decreases with θ. Meanwhile, Algorithm 2 always achieves over 99% of the optimal utility.

VIII. Conclusion

In this paper, we studied the problem of simultaneously assigning threads to servers and allocating server resources to maximize total utility. Each thread was modeled by a concave utility function. We showed that the problem is NP-hard, even when there are only two servers. We also presented two algorithms with approximation ratio $2(\sqrt{2}-1) > 0.828$, running in times $O(mn^2 + n(\log mC)^2)$ and $O(n(\log mC)^2)$, respectively. We tested our algorithms on multiple types of threads, and found that we achieve over 99% of the optimal utility on average. We also perform up to 5.7 times better than several heuristical methods.

In our model we considered homogeneous servers with the same amount of resources and also a single resource type. In the future, we would like to extend our algorithm to accommodate heterogeneous servers with different capacities and multiple types resources. Also, in practice the utility functions of threads may change over time. Thus, we would like to integrate online performance measurements into our algorithms to produce dynamically optimal assignments. Finally, we are interested in applying our methods in real-world systems such as cloud computers and datacenters.

References

[1] V. Srinivasan, T. R. Puzak, and P. G. Emma. Cache miss behavior, is it $\sqrt{2}$. *Proceedings of the 3rd Conference on Computing Frontiers*, 2006

[2] D. Thiebaut. On the fractal dimension of computer programs and its application to the prediction of the cache miss ratio. *IEEE Transactions on Computers*, 1989, vol. 38, no. 7

[3] E. Suh, L. Rudolph, S. Devadas. Dynamic partitioning of shared resource memory. *Journal of Supercomputing Architecture*, 2002

[4] M. K. Qureshi, Y. N. Patt. Utility-based resource partitioning: a low-overhead, high-performance, runtime mechanism to partition shared resources. *IEEE/ACM International Symposium on Microarchitecture*, 2006

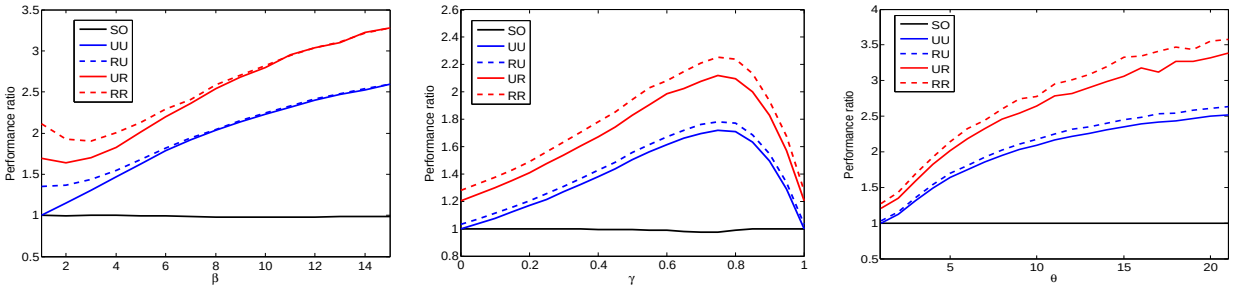

Fig. 3. Performance of Algorithm 2 versus SO,UU,RU,UR, and RR as a function of β, γ and θ under the discrete distribution.

[5] B. Urgaonkar, A. Rosenberg, P. Shenoy. Application placement on a cluster of servers. *International Journal of Foundations of Computer Science*, 2007, vol. 18, no. 05, pp. 1023-1041

[6] A. Karve, T. Kimbrel, G. Pacifici, M. Spreitzer, M. Steinder, M. Sviridenko, and A. Tantawi. Dynamic placement for clustered web applications. *Proceedings of the 15th International Conference on World Wide Web*, 2006

[7] M. Becchi, P. Crowley. Dynamic thread assignment on heterogeneous multiprocessor architectures. *Proceedings of the 3rd Conference on Computing Frontiers*, 2006

[8] P. Radojković, V. Čakarević, M. Moretó, J. Verdú, A. Pajuelo, F. J. Cazorla, M. Nemirovsky, M. Valero. Optimal task assignment in multithreaded processors: a statistical approach. *Architectural Support for Programming Languages and Operating Systems*, 2012

[9] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, R. P. Doyle. Managing energy and server resources in hosting centers. *ACM Symposium on Operating Systems Principles*, 2001

[10] W. Zhao, Z. Wang, Y. Luo. Dynamic memory balancing for virtual machines. *Conference on Virtual Execution Environments*, 2009

[11] P. Lai, R. Fan. Fast optimal nonconcave resource allocation. *IEEE Conference on Computer Communications (INFOCOM)*, 2015

[12] B. Fox. Discrete optimization via marginal analysis. *Management Science*, 1966, vol. 13, pp. 909-918

[13] Y. L. Jiang, X. P. Shen, C. Jie, R. Tripathi. Analysis and approximation of optimal co-scheduling on chip multiprocessors. *Proceedings of International Conference on Parallel Architectures and Compilation Techniques*, 2008

[14] K. Tian, Y. L. Jiang, X. P. Shen, W. Z. Mao. A study on optimally co-scheduling jobs of different lengths on chip multiprocessors. *Proceedings of the 6th ACM Conference on Computing Frontiers*, 2009

[15] S. Zhuravlev, J. C. Saez, S. Blagodurov, A. Fedorova and M. Prieto. Survey of scheduling techniques for addressing shared resources in multicore processors. *ACM Computing Survey*, 2012

[16] Z. Galil. A fast selection algorithm and the problem of optimum distribution of effort. *Journal of the Association for Computing Machinery*, 1979, vol. 26, no. 1, pp. 58-64

[17] H. Kellerer, U. Pfereschy, D. Pisinger. Knapsack Problems. *Springer-Verlag Berlin Heidelberg*, 2004, pp. 329

[18] G.V. Gens and E.V. Levner. An approximate binary search algorithm for the multiple choice knapsack problem. *Information Processing Letters*, 1998, vol. 67, 261-265

[19] E. L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 1979, vol. 4, no. 4, pp. 339-356

[20] A. Neebe and D. Dannenbring. Algorithms for a specialized segregated storage problem. *Technical Report 77-5*, University of North Carolina, 1977

[21] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. *Proceedings of the 11th annual ACM-SIAM Symposium on Discrete Algorithms*, 2000

[22] M. Garey and D. Johnson. Computers and intractability: A guide to the theory of NP-Completeness. *W. H. Freeman & Co.*, 1979

[23] M. Lin, A. Wierman, L.L.H. Andrew, E. Thereska. Dynamic right-sizing for power-proportional data centers. *Proceedings of IEEE INFOCOM*, 2011.